

# HOW IT WORKS

## Firewalls

The internet has made mass communication easy, but it's also opened the door to a host of security problems. **Simon Handby** explains how firewalls protect your PC

**C**omputers, and the connections between them, used to be few and far between, but the internet has made it cheap and easy for people to find and connect to each other. With more online users than ever eager to exchange messages, information, goods and even money, the connections between our computers are changing the way we live.

However, human nature being what it is, a small number of the millions of online users are bound to be up to no good. So just as you'd make sure to lock the doors of a house in the middle of a busy city, you need some way to stop strangers wandering into your private network.

A firewall is a device that protects a computer or computer network from unwanted traffic from untrusted sources. Firewalls were first used in the 1980s, after hackers began to exploit vulnerabilities in the computers and other hardware handling internet and email traffic. With the advent of always-on broadband connections they have become essential for home users, too.

If you have Windows Vista or Windows XP Service Pack 2, a simple but useful software firewall is probably already protecting your computer from most obvious attacks. Ideally, however, you should also have a hardware device that stops internet users from gaining entry to your network in the first place. Fortunately, most modern broadband modem/routers contain a sophisticated firewall that does just that.

### PACKET IN

Described simply, a firewall applies a set of rules to internet traffic, controlling whether it is allowed to pass in to and out from your network, and where it can go when it does. Where a single computer is connected directly to the internet using a dial-up or broadband modem, the modem is 'transparent' to traffic, obediently relaying data to and from the internet to the computer. In such cases, the computer's operating system, or an application running on it, must act as the firewall.

The basis of any firewall is the inspection of the packets of data sent out to and returned from the internet. By comparing information contained in each packet against a list of preconfigured and user-defined rules, the firewall will know whether or not to let each one pass. This technique is known as packet filtering.

Early firewalls operated solely on this basis, judging each packet by criteria such as its source, destination and the type of data it contained. However, it is hard to control data using rigid rules without shutting out valid traffic or accidentally allowing in unwanted traffic. What's more, it's possible to spoof much of the information within a data packet's header, replacing it with false values that could trick a

simple, rule-based filter. Should a hacker find out the details of a valid data source, they could potentially disguise malicious traffic as data of a different type and thus gain entry to the trusted network.

### READY FOR INSPECTION

Newer firewalls evaluate traffic in a more sophisticated way, effectively blocking this exploit. Instead of considering each packet in isolation, they are able to understand that individual packets belong to coherent streams of data, so they can spot interlopers.

Evaluating packets in this way is known as stateful packet inspection (SPI). When a computer on the private network first contacts a public computer, the two undergo a brief handshaking process to establish a connection. At this point, the firewall screens the outgoing request and the reply to ensure that the current set of rules regarding IP addresses and port numbers will permit the conversation.

Provided this is the case, the traffic is allowed and the firewall starts tracking the state of the established connection. It does this by storing important details such as IP addresses, port numbers and the sequence numbers of individual packets in its memory. Packets returned from the internet are compared against the known state of the connection, and any that are identified as not matching (with an identification number substantially out of sequence, for example) are discarded.

SPI firewalls don't require particularly powerful processors, because once a connection is established it's comparatively easy to check whether packets belong to an existing, pre-screened session. However, maintaining the state of a connection requires memory, so idle connections are closed down after a period of time.

An SPI firewall will remove the need for individual rules that allow predefined traffic from the internet. These firewalls can track and permit replies to specific requests originating from the private network, so they are self-configuring for day-to-day inbound traffic.

A third type of firewall seen in certain corporate networks is the proxy or application-layer firewall. This examines network traffic on a higher logical level, and consequently has a broader understanding of how it should behave. Say an attempt was made to sneak in a blocked protocol – such as one typically used in file-sharing applications – on a valid port number such as 80 (used for web traffic). An application layer firewall would be able to spot the mismatch between the port and the protocol's behaviour, and block it accordingly.

Generally speaking, a firewall will block any inbound traffic unless it has been specifically

requested by a computer within your network, which is often referred to as the private or trusted network. However, before a hacker can even try to access your private network, they need to know it's there. By running an automated probe of large numbers of random IP addresses, a hacker can build a list of addresses from which they receive an answer, which confirms that there's some type of equipment connected at that address. Later, they can return and run a comprehensive scan for vulnerabilities.

The most basic function performed by an SPI firewall, but one of the most important, is to ignore any unsolicited inbound traffic from the internet. By doing so, your network is invisible to anyone who is casually probing a large number of IP addresses, and it won't attract any further attention.

### LOST IN TRANSLATION

By pretending to the world that it's not there, a firewall can avoid the attentions of random strangers or automated worms, but anyone who knows your public IP address could target your network specifically. Although this is comparatively unlikely, a firewall needs to be able to withstand a more targeted attack.

One important feature of a typical home firewall, however, recognises the fact that it's generally not the firewall that a hacker is interested in, but the computers that lie behind it. Network address translation (NAT) was originally devised mainly as a way to economise on the limited number of public IP addresses available, but it's a boon for security, too.

When you access the internet from one computer with a modem, or via a gateway that doesn't use NAT, your computer is assigned a public IP address to which any internet-connected device can address information. This is essential for inbound internet traffic to reach it, but it's also a security risk.

When you access the internet via a firewall that uses NAT, however, the firewall is assigned the public IP address, and the computers and other devices connected behind it can each be given an IP address in a private range, such as 192.168.0.XXX. Hackers can't send traffic across the internet to such addresses, and with just one public IP address showing they remain blind to the true layout of your network.

NAT works in real time by replacing the source IP address (192.168.0.7, for example) of outgoing data packets with the public IP address of the router (such as 87.194.0.50). This ensures that replies from a remote computer such as a web server are returned to the firewall's public interface. When packets are returned from the internet, the firewall changes the destination address back to the private IP address of the originating computer, and passes the packet on.

So far, we have looked at the principal ways by which unwanted traffic can be kept out of your network. However, a firewall is useless if it does not also allow the information that you ask for to come in. A NAT firewall knows when it receives the data that a computer on your network requested, and it is able to return it to the right computer.

A typical modem/router's firewall uses a particular type of NAT called port address translation (PAT). It does everything described so far for NAT, but it is designed to cope with multiple computers within the private network. It does this by also replacing the source port number (from the computer that is making the request) contained within the data packet with another port number, and keeping a record of the change.

Although all the data returned by the internet arrives at your single public IP address, it is also addressed with the port number of the requesting device – in this case, the one that your firewall has specified. By comparing this to its records, the firewall knows which machine in the private network asked for the data. It replaces the port number with the one originally used by the computer, and forwards the packet on to it.

UNSOLICITED GOODS

Blocking all unsolicited traffic will help to keep your network safe, but there may be times when you want to allow in selected traffic types. If you run a small business with its own email server, say, then internet mail servers won't be able to deliver to your domain unless you configure your firewall to allow the traffic through.

In most firewalls, this is achieved through the creation of a rule that gives permission to a single type of traffic to enter the trusted network. The aim when configuring such a rule is to make it as restrictive as possible. In the case of email, a typical firewall rule would allow traffic from any public IP address through to only the IP address of your private network's email server. This alone would allow any type of communication, though, so the rule should also restrict the permitted traffic to port 25, which is used by the simple mail transfer protocol (SMTP).

Allowing traffic into your network is a risk, however. If your mail server's software has an unpatched vulnerability, for example, it might be possible for a hacker to slip through the firewall

FURTHER INFORMATION

Wikipedia entry on firewalls  
[www.wikipedia.org/wiki/firewall](http://www.wikipedia.org/wiki/firewall)  
Detailed Cisco guide to firewall types  
[www.tinyurl.com/yqqdb](http://www.tinyurl.com/yqqdb)  
Wikipedia entry on Internet Protocol version 4 (IPv4)  
[www.wikipedia.org/wiki/IPv4](http://www.wikipedia.org/wiki/IPv4)


and gain control of your server. This might be bad enough, but if you're really unlucky, they might be able to access other computers on your network from there.

One way that enterprise-class firewalls remove this risk is by having a third network interface that allows the creation of a demilitarised zone (DMZ), separate to both the internet and the private network. Computers that need to be accessible from the internet, such as email, web and FTP servers, are placed in this network. They are protected by the firewall, and the administrator can create rules that allow the minimum necessary traffic to them.

The firewall protects the DMZ from the internet, and also protects the private network from the DMZ, so any hacker who compromises a server in the DMZ still has to find a route through the firewall before reaching the private network.

Recently, as the firewalls in modem/routers have become more sophisticated, many have included a DMZ feature. However, this often exposes the private IP address to inbound internet traffic, which is a security risk. ☹

Building a wall How a firewall protects a network



The firewall's job is to protect the internal network from internet threats. It's a two-way process. Traffic is filtered on the way out of the network, to stop internal computers from accessing services they shouldn't, and filtered on the way in to stop hackers from attacking. A second form of defence is network address translation (NAT). The

firewall uses an external public IP address for communication, but hides the internal network behind an internal IP address. This address is not internet-visible, making it harder for hackers to gain access. Finally, the DMZ is used to separate certain computers, such as servers, from the internal network. If the server is hacked, the rest of the network is safe.

DIAGRAM: MIKE HARDING